

LSI 01-524/1C

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR PATENT

ON

ARCHITECTURE FOR A SEA OF PLATFORMS

BY

CHRISTOPHER HAMLIN
310 JENSEN SPRINGS RD
LOS GATOS, CA 95033
CITIZEN OF USA

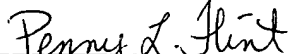
CERTIFICATE OF MAILING BY "EXPRESS MAIL"

"Express Mail" Mailing Label Number: EV 338 284 132 US

Date of Deposit: July 23, 2003

I hereby certify that this correspondence is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. § 1.10 on the date indicated above and is addressed to MS Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450

BY:



Penny L. Flint

ARCHITECTURE FOR A SEA OF PLATFORMS

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] The present application is a continuation-in-part of U.S. Patent Application No. 10/044,781, filed January 10, 2002, now pending. The present application is also a continuation-in-part of U.S. Patent Application No. 10/135,189, filed April 30, 2002, now pending.

[0002] The present application hereby incorporates the following United States Patent Applications by reference in their entirety:

<u>Attorney Docket Number</u>	<u>U.S. Patent Application No.</u>	<u>Filing Date</u>
LSI 01-390	10/015,194	November 20, 2001
LSI 01-488	10/021,414	October 30, 2001
LSI 01-489	10/021,619	October 30, 2001
LSI 01-490	10/021,696	October 30, 2001
LSI 01-524B	10/044,781	January 10, 2002
LSI 01-543	10/135,189	April 30, 2002
LSI 00-695	09/842,335	April 25, 2001
LSI 01-827	10/034,839	December 27, 2001
LSI 01-828	10/034,648	December 27, 2001

FIELD OF THE INVENTION

[0003] The present invention generally relates to the field of semiconductors and semiconductor design, and particularly to an architecture for a configurable platform.

BACKGROUND OF THE INVENTION

[0004] Integrated circuits have become a necessary part of everyday modern society. From wireless phones and information handling systems, to household appliances and data storage systems, a wide range of integrated circuits are utilized to provide a broad

range of functionality. To provide this functionality, integrated circuits may need to be specialized to have the functions necessary to achieve the desired results, such as through the provision of an application specific integrated circuit (ASIC). An ASIC is typically optimized for a given function set, thereby enabling the circuit to perform the functions in an optimized manner. However, there may be a wide variety of end-users desiring such targeted functionality, with each user desiring different functionality for different uses.

[0005] Additionally, more and more functions are being included within each integrated circuit. While providing a semiconductor device that includes a greater range of functions supported by the device, inclusion of this range further complicates the design and increases the complexity of the manufacturing process. Further, such targeted functionality may render the device suitable for a narrow range of consumers, thereby at least partially removing an “economy of scale” effect that may be realized by selling greater quantities of the device.

[0006] Thus, the application specific integrated circuit business is confronted by the contradiction that the costs of design and manufacture dictate high volumes of complex designs. Because of this, the number of companies fielding such custom designs is dwindling in the face of those rapidly escalating costs.

[0007] Therefore, it would be desirable to provide a platform architecture of the present invention.

SUMMARY OF THE INVENTION

[0008] Accordingly, the present invention is directed to platform architecture used for integrated circuit design. In a first aspect of the present invention, a system for providing distributed dynamic functionality in an electronic environment includes a plurality of platforms. The platforms are suitable for providing a logic function, and include embedded programmable logic, memory and a reconfigurable core. The logic, memory

and reconfigurable core are communicatively coupled via a fabric interconnect. A map is also included which expresses logic functions of the plurality of platforms.

[0009] In a second aspect of the present invention, a method for providing an executable suitable for being employed by a plurality of platforms includes receiving a program of instructions and determining availability of a plurality of platforms for performing the program of instructions. The plurality of platforms include embedded programmable logic, memory and a reconfigurable core. The logic, memory and reconfigurable core are communicatively coupled via a fabric interconnect that is isochronous. The availability of the platforms includes at least one of load value of the platforms and functionality of the platforms. The program of instructions is translated into an executable suitable for operation by the plurality of platforms based on the determined availability.

[0010] In a third aspect of the present invention, a system for providing distributed dynamic functionality in an electronic environment includes a plurality of platforms communicatively coupled via an isochronous fabric. The platforms are suitable for providing a logic function, and include embedded programmable logic, memory and a reconfigurable core, which are communicatively coupled. A map is also included which expresses availability of the plurality of platforms for performing a logic function.

[0011] It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the invention as claimed. The accompanying drawings, which are incorporated in and constitute a part of the specification, illustrate an embodiment of the invention and together with the general description, serve to explain the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] The numerous advantages of the present invention may be better understood by those skilled in the art by reference to the accompanying figures in which:

[0013] FIG. 1 is an illustration of an embodiment of the present invention wherein a platform operable to embody the present invention is shown; and

[0014] FIG. 2 is an illustration of an embodiment of the present invention wherein a plurality of platforms as shown in FIG. 1 are provided in a fabric.

DETAILED DESCRIPTION OF THE INVENTION

[0015] Reference will now be made in detail to the presently preferred embodiments of the invention, examples of which are illustrated in the accompanying drawings.

[0016] Referring generally now to FIGS. 1 through 2, exemplary embodiments of the present invention are shown. The present invention provides an architecture for an element for use as a design component in application specific integrated circuit (ASIC) or semiconductor design. Technical questions involving complex device design may be thought of in two general areas, metamethodology and platform architecture. Metamethodology is a formal organizing architecture for defining and managing arbitrary semiconductor design flows in predictable, efficient ways, which may be tailored to specific product and process characteristics. For instance, different flows may have different combinations of tools in them, and thus, the successful operations may require the imposition of automated rule-based assessment of progress and error in order to allow designers to work productively at a higher level of abstraction than is typical today.

[0017] As semiconductor design progresses along the deep sub-micron roadmap, and as the challenge of managing complexity becomes ever greater, the need to define building blocks of designs at higher levels of abstraction becomes even more pressing. Such a higher-level building block includes a “platform.”

[0018] Referring now to FIG. 1, an embodiment 100 of the present invention is shown wherein elements of an exemplary platform are provided. A platform may include a combination of the following elements: (1) embedded programmable logic 102, which in some contemplated embodiments is analogous to field programmable gate array (FPGA) or complex programmable logic device (CPLD) cores that FPGA companies sell as complete devices; (2) reconfigurable cores 104 such as a fundamental processor element to which instruction-specialized, application-specific instruction set extensions may be added; (3) an advanced interconnect 106, which in contemplated embodiments is scalable, and may be isochronous; (4) software models and heuristics 108; and (5) specialized memories 110, which may include nonvolatile structures like MRAM, which is a memory that is based on the magneto resistive effect, as well as other memories as contemplated by a person of ordinary skill in the art. Specialized one-time programmable flash memory may also be included.

[0019] Programmable logic 102 components may include blocks of programmable gate arrays, “seas of adders”, CPLD structures, and other suitable programmable circuit elements definable from a stored representation, either at power-on, dynamically while in operation, and the like.

[0020] Reconfigurable cores 104 may include a base processor design, plus instruction set extensions designed to carry out function-specific logical and arithmetic operations with optimal efficiency. For instance, such reconfigurable cores may implement digital signal processing instruction set enhancements.

[0021] An interconnect 106 architecture is provided to allow the programmable logic and the reconfigurable cores to communicate with one another and with associated memory blocks. Such an architecture may define a transport which is scalable in bandwidth and is inherently isochronous. Further, it may be realizable within a switching fabric, which permits complex, adaptive, interconnect and access paths to be defined on the fly.

Isochrony, through a universal time base, may simplify the problems of closing timing in complex designs.

[0022] The software for the platform may include the development environment and its interface to the metamethodology, as well as software IP cores, which may be implemented on the platform components, and the customer-developed custom code implementing proprietary or functionally specific routines.

[0023] There are numerous ways of structuring logic blocks: such as the number of logic elements in an organization; the kinds of cores and the numbers of cores; whether or not DSP-specific characteristics are included, and the like. Additionally, the characteristics of the interconnect, software, and memory may all be varied in innumerable different ways without departing from the spirit and scope of the present invention.

[0024] When a collection of electronic resources including digital circuits, processors, and analog circuits is integrated by means of an interconnection network with regular or partially regular structure it is termed a fabric. A particular type of fabric in which each node comprises an instance of a “platform” is termed a “sea of platforms”.

[0025] Of particular importance in a sea of platforms is the signaling discipline whereby the elements in the fabric are clocked and are coordinated in time. In traditional practice, for instance, a regular digital clock pulse is supplied to each element, and the edges of the clock(s) are aligned locally using PLLs or DLLs, for example.

[0026] In a sea of platforms architecture the physical resources are abstracted to a considerable extent from the explicit function(s) they will be called upon to implement. In the same fashion it is desirable to abstract as far as possible the timing resources which the nodes in the sea of platforms fabric rely upon, in the interests of efficiency and reliability.

[0027] In the sea of platforms architecture, the present invention proposes that this be done by implementing a plesiochronous signaling discipline by means of which every interconnected element in the sea of platforms is provided with just the timing resources necessary to accomplish its task, but no more. In other words, excess accuracy or precision in the provision of timing resources to a node or unit in the fabric is undesirable because it allocates energy or physical components unnecessarily and therefore inefficiently to the instantaneous timing requirements of the unit.

[0028] Under the present invention, a protocol is used to establish a pre-defined "local regime" for timing purposes, in terms of which the node in question will conduct its operations. It may vary with respect to the activities being conducted simultaneously in other nodes, to explicit degrees which are asserted by the protocol. At any time, the node may request or may be required to establish a new degree of coordination or synchronization with other members of the fabric, under the control of the protocol conveyed to the nodes in question.

[0029] Such signaling, which differs, for example, from jitter management schemes, falls broadly under the classification of plesiochronous architectures. Its central advantage in the sea of platforms case is that timing may be permitted to vary freely with respect to precision and accuracy where appropriate within the fabric, but may at any point be coordinated or synchronized with external entities whose state has arrived at a requirement for the re-imposition of a greater degree of e.g. accuracy or precision.

[0030] Such a plesiochronous sea of platforms depends upon having access to a universal timebase and set of clocking resources which can be sampled and parceled out via the protocol as needed to the participating nodes; the specific uses of the protocol may be programmatic or algorithmic, and can answer both deterministic and stochastic process requirements as they arise within the sea of platforms.

[0031] The present invention provides a platform, which may be embedded in a methodological framework that allows designers to work with the platform at a high level of abstraction. The complexity of the interactions between elements may be so great that by providing abstractions, the interactions may be rendered manageable and tractable for designers that will use them practically.

[0032] Referring now to FIG. 2, an embodiment 200 of the present invention is shown wherein an architecture is configured as a sea of platforms. In this embodiment, a fabric is shown including a plurality of platforms. Preferably, the platforms are connected utilizing similar protocols, interconnect technology and interconnect architecture to unify the platforms in a single fabric.

[0033] Thus, a structure of the resources may be provided including the memory, reconfigurable cores, embedded programmable blocks, software and interconnect, which communicates and intercommunicates coherently on an isochronous fabric. Such a structure may be suitable for providing a “programmable ASIC”. For instance, an application-neutral device may be constructed with potential for accepting complex logic and definitions that are programmed completely independently of the fabrication of the device. This may be thought of as a decoupled model, in which the contemplated implementation of the device as used by an end-user is decoupled from the process of designing and manufacturing the physical device itself.

[0034] Each of these intersections on this fabric, may contain one or more processors, embedded programmable logic, memory, software capabilities and its own interconnect internally, as described in relation to FIG. 1. By providing this fabric including these reconfigurable cores, the present invention may provide a targeted semiconductor providing desired functionality without requiring specialized design and manufacturing processes as previously required in ASIC manufacture.

[0035] For example, a manufacturer may receive a register transfer level (RTL) definition of a solution to a problem from a customer for a specification. A customer may want to create, for instance, a communications device, a storage device, a controller, a switching product, a game controller for consumer application, a satellite TV set-up box, and the like, and supply an RTL specification for the desired device to the manufacturer.

[0036] By utilizing the present invention, the RTL may be mapped into the platform of the present invention to provide the desired functionality as indicated by the RTL. For instance, instruction set architecture extensions may be utilized for mapping to the reconfigurable core the desired functionality. The extensions may be crafted to solve efficiently and specifically problems in encryption, in encoding or decoding, in modulation, in signal processing, in data transformation of various kinds, and the like. Additionally, abstract logic functions may be implemented, such as specialized shift registers, multiplexers (MUX), and the like as contemplated by a person of ordinary skill in the art. Thus, an instruction set extension may be affiliated with a wide range of functionality.

[0037] An embedded programmable logic core (EPLC) block may be tied to a set of instruction set extensions such that the EPLC block would have responsibility under software control for invoking any of several extensions to be active in a particular temporal episode. In this way, the “personality” of a reconfigurable core may be changed dynamically under the control of this EPLC mechanism. Thus, in an aspect of the present invention, this may enable the mapping of an RTL efficiently into a sea of platforms. The constituents of one platform may allow an EPLC block to play a role when choosing appropriate instruction set extensions as needed, given the temporal evolution of the function that the block is fulfilling.

ISOCHRONOUS FUNCTIONALITY

[0038] Referring again to FIG. 2, by providing an isochronous fabric, the register files of multiple processors in multiple platforms may be utilized as dynamically extensible. In other words, it is possible because of the isochronous characteristic of the fabric, without additional software or additional overhead, to synchronize and coordinate the instruction-set extension operations on these register files over as many platforms as needed, which may be thought of as extending horizontally across the fabric as desired to achieve the necessary resources, such as processor power and the like, to fulfill a particular complex logic function.

[0039] Arbitrary sets of logic functions may be deployed across register files, and the instruction set extensions treated as general logic engines that are reconfigurable “on-the-fly,” on a cycle-by-cycle basis. For instance, given, (a), the provision of a proper instruction set extension or extensions to coordinate discrete, distinctive, different instruction set extensions on a cycle-by-cycle basis; and (b), that the execution is synchronized, for instance, to ensure that the right instruction set extension is invoked in the right set of reconfigurable cores at the right cycle, the functions may be deployed across platforms, operating as logic engines, as needed. Thus, by knowing the functionality of a register, where that functionality is located, and the function of a register at a given point in time, logic functions may be targeted to provide the functionality.

[0040] For example, in terms of actual behavior of applications in the real world, loads vary, and functions vary as loads vary. By utilizing the present invention, spatial distribution of functions across the array may vary as a function of the dynamic changes in the functional load that is actually being asked of a particular device.

[0041] To track these changes and provide the functionality, a map may be maintained indicating the functionality of the platforms. For instance, in an aspect of the present invention, a master map is maintained of the instantaneous distribution of functions across the platforms. Such a map may be thought of as a functional virtualization, in which the map indicating corresponding functions and locations is fully virtualized. Thus, functional virtualization may be provided in addition to a general logic capability previously discussed.

[0042] Because the isochronous foundation of this embodiment of the present invention, the isochronous fabric provides coordination synchronization without the bookkeeping or overhead which may be associated utilizing other methods. By providing a mapping of particular components, i.e. what the particular components are set up to do what particular component of a function at what point in time, desired functionality may be achieved in a coordinated fashion.

COMPILER

[0043] In an additional aspect of the present invention, a smart compiler is provided which “understands” how to manage and develop a binary executable for a particular instruction set extension. Further, the compiler technology may be generalized so that it has the property of extending this understanding, so that it may track which extensions are mapped to which particular set of processors, and understands temporally the load value, i.e. the cycle by cycle availability of a resource of a particular kind.

[0044] In effect, the compiler technology extends horizontally across processor function sets, so that the compiler, when an application, methodology and like program of instructions is expressed and translated to the compiler, the compiler may determine availability of the resources. Additionally, through the use of an isochronous fabric, there is no overhead associated with altering the connections. Reconfiguring the functionality of a device employing the platforms may be accomplished through changing the map.

[0045] The compiler technology of the present invention may implement this space/time view through an arbitrarily extensible very large instruction word (VLIW) architecture that is variable. For example, although the architecture has been used in multimedia engines, the width is cycle-by-cycle variable according to this aspect of the present invention.

[0046] The advantages of such a “smart” compiler are numerous. For instance, in cache management, a variety of considerations may be accounted for, such as latency and the resultant performance penalty, associated overhead of flushing the cache versus maintaining a function in place, and the like. Thus, a compiler of the present invention may optimize operations performed by the platforms.

[0047] Therefore, in an embodiment of the present invention, a compiler is provided that is capable of maintaining space/time mapping of the instruction set extensions over an isochronous fabric so that cycle-by-cycle ability is maintained to affiliate objects and communicatively couple them as desired.

MAPPING

[0048] To coordinate and provide desired functionality, a map of the present invention is provided. For instance, in an aspect of the present invention, RTL is expressed, in terms of combinations of instruction set architecture extensions and embedded programmable logic core (EPLC) blocks.

[0049] For example, in an aspect of the present invention, a map is provided for describing functions of platforms of the present invention expressed in a graph-theoretic manner. A map may be provided as a graph, for instance, employing graph coloring and

with efficient graph-traversal algorithms to describe the interaction and functionality of the components.

[0050] Formalisms may be employed for expressing functions, such as MUXs, latches, codecs, and like logic functions and re-expressing the functions in terms of efficient instruction set architecture (ISA) extensions. Preferably, the extensions take into account that some components of the ISA may be modified on a cycle-by-cycle basis on the one hand and may be varied in width on the other hand.

[0051] For instance, standard library functions, such as concrete, practical, standard functions in an ASIC, may be expressed as mathematical abstracts. These mathematical abstracts may be expressed as instruction set extensions and EPLC adjuncts that will allow these instructions to be manipulated rapidly. In one contemplated embodiment, ASIC library functions are implemented with minimal overhead penalties and space penalties.

[0052] In this way, the present invention provides two degrees of freedom in finding optimal expressions of logic functions that can be deployed across the sea of platforms of the present invention.

RESOURCE FABRIC AND BROKEN SYMMETRY

[0053] The use of a switched fabric to implement a “sea-of-platforms” architecture has been described above. This section describes a “resource fabric”, which might preferably be based upon a “sea-of-platforms” such as described above, and whose preferred embodiment entails the use of “broken symmetry” in the optimization of the fabric characteristics.

[0054] Semiconductor devices may be thought of as converting collections of transistors into functions. This may be achieved by ordering logic operations and state transitions in time in a manner which accomplishes a sought-after operation on data. As it has become possible to manufacture devices with more and more transistors on them, the functions implemented have become more complex and diverse. This evolution has been accompanied by an attendant increase in the complexity of the pattern of connections between groups of transistors on devices, and also by a tendency to group larger and larger sets of transistors according to abstractions operating at higher and higher levels. Thus functions become available to designers in abstracted forms which make it unnecessary to know how a particular grouping of transistors is organized to implement the function.

[0055] As devices approach the one billion transistor level of complexity, it is desirable to accompany this growth with a suitable extension of the level of abstraction employed to efficiently gather together groups of transistors which accomplish a useful, abstracted purpose which designers can readily incorporate into larger design structures.

[0056] To this end, the present invention pertains to a new level of abstraction in terms of which groupings of transistors on a semiconductor device may be organized, which extends the level of abstraction from a “function” to a “resource”. Furthermore, it details the implementation of such abstracted resources within a “resource fabric” which might be realized in a number of ways, but whose preferred manner of implementation is a fabric characterized by “broken symmetry” in any of the dimensions of a) physical 3-dimensional space; b) temporal space; or c) code space. These concepts are further elaborated below.

[0057] The economics of manufacturing very high complexity devices increasingly favor the reuse of physical assets on a die wherever possible, because the fixed costs of factors of production such as mask sets is rising nearly exponentially in some cases.

[0058] As a result, just as it proved fruitful in standard cell ASIC design to abstract sets of transistors into functional groups implementing specific tasks such as decoding or deserialization or filtering, the present invention contemplates the further abstraction of sets of transistors and sub-functions into “resources”, which may be combined in specific ways to achieve certain functional objectives.

[0059] As an illustration of the further abstraction of sets of transistors into “resources” (though not an exhaustive one) the elements of a high complexity device could be thought of as falling into any of three categories: Transformative, Coordinating, and State Management.

[0060] Notice that none of these three categories is an end function proper; rather these are abstracted descriptions of resource types needed by various different functional tasks. Under the heading Transformative Resources might fall Conversion and Translation of data, Modulation, Calculation, Coding and Decoding, and so on. Under the heading of Coordinating Resources might fall Timing, Synchronization, Buffering, Caching, and so on. State Management Resources might include control of non-volatile storage structures, tree traversal capabilities, and so on.

[0061] The elevation of the level of abstraction applied to sets of resources in this manner requires a suitably flexible, malleable strategy for interconnecting the resources so that they may be fashioned into the arrangements which best satisfy an ultimate functional objective. The structure described in the aforementioned “sea-of-platforms” disclosure is appropriate for this purpose, if extended in the manner now to be described.

[0062] One of the goals of device architecture is to maximize the efficiency with which a general semiconductor organization which minimizes manufacturing cost can translate a specific customer design into product form. Device organizations span a continuum from

extremely specific, which are suitable for very high volume, cost sensitive applications, to very general (illustrated by FPGAs) which match cost insensitive low volume markets.

[0063] The ideal device architecture is one in which a very high degree of generality is sustained (thereby lowering manufacturing and inventory cost and risk) without sacrificing efficiency in accomodating a large range of specific product designs.

[0064] Hence the present invention employs the sea-of-platforms as a vehicle for unifying a flexible and malleable collection of abstracted resources in such a way as to optimize their use for specific designs.

[0065] A progression is envisaged whereby the fabric structures evolve over time (perhaps through the 65 nm node), according to the following sequence:

- 1) regular matrix, asynchronous signaling;
- 2) regular matrix, isochronous signaling;
- 3) 'perturbed' matrix, plesiochronous signaling;
- 4) multidimensional symmetry-breaking matrix (encompassing the dimensions of physical space, temporality, and code space)

[0066] In each instance the fabric represents a set of interconnects between nodes, each node comprising a set of resources. The important feature of the present invention is that the need for efficiency in mapping between a very general organization and very specific applications is satisfied because the sea-of-platforms set of resources is freed from fixity in physical placement and from fixity with respect to time, in the sense that it becomes possible to schedule consumption of resources throughout the fabric, and to incorporate such schedules into the design environment which is used to translate specific designs into the resource fabric.

[0067] The efficiency of this mapping ultimately resolves to selection of a device family member with just sufficient spatial and temporal resources to meet the requirements of a design; the fit between the specific design and the fabric-based resource set is optimized through the placement and scheduling operations.

[0068] In effect this efficiency optimization may be thought of as an extension of conventional place-and-route operations, which are applied in two-dimensional space to the physical matrix of transistors on a device, to a higher dimensional abstract space in which time is also available for “placement”, except that in the case of time the resource consumption takes the form of scheduling. But from an abstract point of view the result is the same: an asset is consumed, or “placed”, with respect to all other assets in the device, in satisfaction of the design’s specific requirements.

[0069] The problem of inferring that structure which will optimally map a general set of resources to a specific product design with minimal inefficiency is of the sort called “inverse”. The present invention applies a property called “broken symmetry” to this inverse problem. Broken symmetry arises in many contexts with similar properties, in that neither perfectly regular structures such as rectilinear Cartesian coordinate systems, nor highly irregular structures such as the random-seeming patterns emerging from genetic programming solutions, offer the optimal balance for reconciling the specific with the general. Broken symmetry operates to balance the competing demands of specific designs and general resource sets on the physical and temporal resources of a device’s architecture.

[0070] The present invention provides use of an isochronous fabric to tie together and coordinate the use of platform resources, such as through the use of a map, compiler, and the like as described in embodiments previously mentioned. It is believed that the platform architecture of the present invention and many of its attendant advantages will be understood by the forgoing description. It is also believed that it will be apparent that

various changes may be made in the form, construction and arrangement of the components thereof without departing from the scope and spirit of the invention or without sacrificing all of its material advantages. The form herein before described being merely an explanatory embodiment thereof, it is the intention of the following claims to encompass and include such changes.